

AFRL-RI-RS-TR-2009-139
Final Technical Report
May 2009



THE GOAL-MEANS (GM) META MODEL

National Security Innovation, Inc.

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

**AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the 88th ABW, Wright-Patterson AFB Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-RI-RS-TR-2009-139 HAS BEEN REVIEWED AND IS APPROVED FOR
PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION
STATEMENT.

FOR THE DIRECTOR:

/s/

/s/

JAMES M. NAGY
Work Unit Manager

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY) MAY 09	2. REPORT TYPE Final	3. DATES COVERED (From - To) Nov 08 – Jan 09		
4. TITLE AND SUBTITLE THE GOAL-MEANS (GM) META MODEL		5a. CONTRACT NUMBER FA8750-08-C-0208		
		5b. GRANT NUMBER N/A		
		5c. PROGRAM ELEMENT NUMBER N/A		
6. AUTHOR(S) Elisa Bienenstock, Pamela Toman, Kari Kelton, Allison Astorino-Courtois, Stacy Lovell-Pfautz, and Robert Popp		5d. PROJECT NUMBER PAIN		
		5e. TASK NUMBER 00		
		5f. WORK UNIT NUMBER 15		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) National Security Innovations, Inc. 8 Faneuil Hall Marketplace Boston, MA 02109		8. PERFORMING ORGANIZATION REPORT NUMBER N/A		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/RIED 525 Brooks Rd. Rome NY 13441-4505		10. SPONSOR/MONITOR'S ACRONYM(S) N/A		
		11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-RI-RS-TR-2009-139		
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 88ABW-2009-2057				
13. SUPPLEMENTARY NOTES				
14. ABSTRACT National Security Innovations (NSI) developed a Goals-Means (GM) Meta-model of leadership decision-making as a component of the IARPA PAINT system. The GM Meta Model instantiates social psychological theories on human decision-making into a java program that provides analysts with a tool to replicate and deconstruct the complex decision calculus. The tool incorporates a range of decision criteria, heuristics and styles to approximate real human decision making to capture the values and priorities of leaders from varied cultural and historical contexts. In addition to modeling individual decision process, the GM meta-model models the effect of group dynamics and social influence to allow simultaneous analysis of multiple decisions from a dynamic set of decision makers.				
15. SUBJECT TERMS Decision models, Goals Means meta-model, Data integrator module, Dissonance, Social influence, Leadership decisions, heuristic, Expected utility, game theory, PAINT				
16. SECURITY CLASSIFICATION OF: a. REPORT U		17. LIMITATION OF ABSTRACT UU	18. NUMBER OF PAGES 45	19a. NAME OF RESPONSIBLE PERSON James M. Nagy
				19b. TELEPHONE NUMBER (Include area code) N/A

TABLE OF CONTENTS

1. EXECUTIVE SUMMARY.....	1
2. INTRODUCTION	3
3. METHODS, ASSUMPTIONS, AND PROCEDURES	6
3.1 RESEARCH AND DEVELOPMENT.....	7
3.1.1 <i>Methods</i>	7
3.1.2 <i>Assumptions</i>	8
3.1.3 <i>Procedures</i>.....	10
3.2 COLLABORATION AND DEMONSTRATIONS	22
3.2.1 <i>Methods</i>	22
3.2.2 <i>Assumptions</i>	22
3.2.3 <i>Procedures</i>.....	24
4. RESULTS AND DISCUSSION	30
4.1 RESEARCH AND DEVELOPMENT.....	30
4.2 COLLABORATION AND DEMONSTRATIONS	31
5. CONCLUSIONS	33
6. RECOMMENDATIONS	34
7. REFERENCES	35
8. LIST(S) OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS.....	38

LIST OF FIGURES

Figure 1: The GM Meta Model, its components and interface to other PAINT modules	1
Figure 2: Gathering data about decision choices unrelated to a question, it is possible to see inside a "black box" and determine which path decision makers are following.....	3
Figure 3: The MADM module assumes that decisions are affected by: 1) the social network, 2) the environmental context, and 3) the preferences of a decision maker	8
Figure 4: Decisions are not made in isolation; aggregation of group opinions must be addressed as well	9
Figure 5: The MADM module addresses a number of questions in turn through its process and method of determining potential final decisions.....	10
Figure 6: Decision matrices formalize the decision-making process	11
Figure 7: Decision matrices can be operated upon using an expected utility/rational choice rule.....	12
Figure 8: Decision matrices can be operated upon using an elimination-by-aspects heuristic.	14
Figure 9: Decision matrices can be operated upon using a lexicographic choice heuristic.....	15
Figure 10: Decision matrices can be operated upon using a maximin choice heuristic	16
Figure 11: Social influence from peers and superiors pushes and pulls actors into new opinions and therefore new decisions.....	20
Figure 12: Accurate data for the decision model is assumed to come from many different performers.....	23
Figure 13: In the integrated model, when a pathway is starved, decision makers respond to fund it more extensively.....	31

LIST OF TABLES

Table 1: The Vroom-Yetton model differentiates three main types of decision makers: autocratic, consultative, and group.....	21
Table 2: An example decision-making data file (.dmd file) format.....	24
Table 3: The MADM module accepts formalized input during instantiation that defines the decision makers, their preferences, and the decisions that may need to be made, in a particular format.....	25
Table 4: The probe to "toggle off" an interest's relevance to individual decision makers can be taken against the decision model	28

ACKNOWLEDGEMENTS

The NSI team would like to thank the following organizations and personnel for their considerable support, commitment and leadership all of which enhanced the quality of our work significantly. First, special thanks to the government sponsors at DNI/IARPA and AFRL, in particular, Peter Brooks, Jim Nagy, Dan Daskiewich and Nancy Roberts. In addition, we would also like to thank the SAIC team of Ted Senator and Lars Hanson. Special thanks to the BAE team of Ed Waltz, JP Parker and Matt Humphreys. And finally, we would also like to thank the Booz Allen Hamilton and MIT teams, in particular, Dan Horn and Paul Barsotti of Booz Allen Hamilton, and Daniel Goldsmith, Michael Siegel and Stuart Madnick of MIT.

1. EXECUTIVE SUMMARY

The Pro-Active Intelligence (PAINT) program sought to develop models and methods to assist the United States Government in revealing the motivations, intentions and activities of individuals and groups that threaten its security. To achieve this goal NSI developed the Goals-Means (GM) Meta Model of leadership decision making. As a component of the PAINT system, the GM Meta Model reveals the true state of a system by modeling decision processes under a range of initial conditions that provides a generalizable and scalable method for forecasting decision outcomes and for reverse-engineering the key factors that differentiate between decision outcomes.

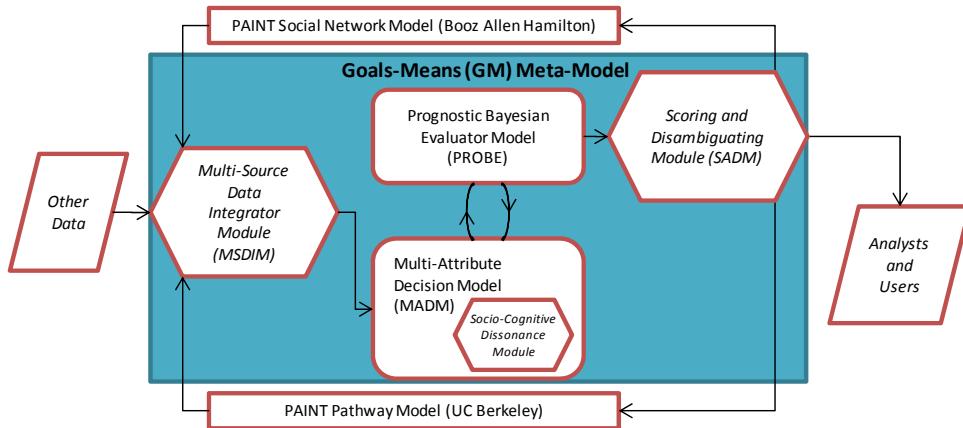


Figure 1: The GM Meta Model, its components and interface to other PAINT modules

The GM Meta Model is a system of four modules (see Figure 1); when integrated, it provides a parameterized analysis of leadership decision making. The primary benefits of the GM Meta Model are six fold:

1. it allows insight into the likely intentions, priorities, interests, strategies and choices of decision makers;
2. it enables comparison of sets of potential decision outcomes as explicitly linked to empirical data analyses;
3. it allows perturbations, modifications and updating of original input data;
4. it reveals the link between assumptions about the true state of a system and sets of potential decision outcomes;
5. it enables tests of the robustness of potential outcomes and identification of probes (subtle changes in inputs that produce differentiating decisions) through sensitivity analysis on the assumptions; and
6. It assists in identifying probes that reveal the state of the world when integrated with other modules in the PAINT system.

The core of the Goals-Means Model is the Multi-Attribute Decision Model (MADM) module which reconstructs the subjective decision-making calculus of actors to model leadership decision making. MADM was designed for integration into the PAINT architecture as a general, scalable model of operational decision making. NSI prioritized the development of the MADM module and integration with other modules in the BAE PAINT architecture to meet the PAINT condensed schedule. Model development ensured that MADM successfully replicated human decision making by prioritizing the following four principles:

1. actors are not always perfectly rational;
2. decisions are influenced by both internal (mental) and external (social) structures;
3. actors are distinct from the roles they fill; and
4. Group behavior is not the sum of its parts.

As a result, the MADM module includes an array of decision-making algorithms including a multi-attribute game theoretic model and a set of heuristic options to model the gamut of decision-making behaviors. In addition the model heavily integrated with the social network models. An important component of this module is the *Socio-Cognitive Dissonance Module (SCDM)* which models the effect of social structures and social interactions on actor decisions.

The results of internal validation tests and the integrated demonstration show that MADM successfully models the decision-making process and appropriately responds to social and organizational environment.

A condensed PAINT schedule required NSI to limit its development focus to component modules required to meet the PAINT demonstration schedule. NSI has produced domain- and platform-independent software components to model multi-actor decision processes, update actor preferences, and invoke multiple decision-making heuristics, with a framework of Java software that can be run stand-alone or integrated into the other PAINT performer modules through the BAE backplane. The development of the *Multi-Source Data Integrator Module (MSDIM)* (to prepare data from varied sources including SME input, other performers models, structured data bases or social network analysis), the *Prognostic Bayesian Evaluator (PROBE)* (to enable sensitivity analysis to parameterize outcomes and access reliability), and the *Scoring and Disambiguating Module (SADM)* (to deconstruct model outputs to identify decisions that reveal important information) were not completed to meet the demonstration requirements of the integrator BAE.

2. INTRODUCTION

The vision of the PAINT program was to develop systematic approaches to reveal what is not readily knowable about the intentions, capabilities and technologies available to foreign entities that threaten US interests. Unlike similar efforts, PAINT was not an attempt to predict future outcomes; instead, its focus was to reveal direct evidence of programs to develop a threatening product or capability, or to reveal indirect evidence or important information to validate, update or support existing analytic models.

NSI supported the PAINT program through contributing: (1) a hybrid leadership decision-making GM Meta Model and (2) meeting demonstration requirements, as led by BAE. The GM Meta Model combines knowable (or assumed, if knowing is impossible) information with a computational model of decision making, in order to reveal aspects of the world not apparent or measurable.

In its original conception, the GM Meta Model would not predict what decision leaders will choose, but rather to depict the range of likely decision outcomes associated with various (personality and context) inputs. The results could be investigated by analysts by tweaking the decision space, thereby exploring the effect of different framings of reality; it could also be explored by machines to parameterize and generate error bars around the potential decisions that a decision maker might make giving the analyst a better sense of the probability distribution function between different potential outcomes.

The vision for the GM Meta Model was to use models of leadership decision-making to reveal aspects of the world that are not readily knowable. Modeling how decision outcomes differ under different initial conditions provides a means of revealing the true state of those initial conditions. While the foreign entity may consciously conceal the initial conditions, these initial conditions may influence many ancillary decisions that are not concealed. Observing these outcomes then provides insight into what is hidden. In this way what is not apparent or measurable can be considered to be a “black box” of information that we do not know *a priori* but nevertheless have an interest in knowing. The insights provided by the GM Meta Model is that through knowledge of (or even fair approximations of) information that *can* be collected from outside that box, it is possible to characterize the contents of the box – that is, the previously-hidden state of the world (see Figure 2).

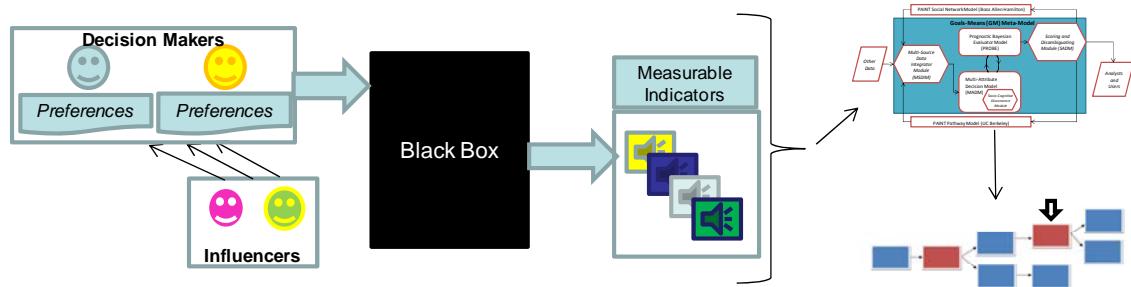


Figure 2: Gathering data about decision choices unrelated to a question, it is possible to see inside a "black box" and determine which path decision makers are following

In the formalization explored by NSI, the information needed to characterize the black box includes who potential decision makers are, what their preferences are, who the people are who influence each of them (that is, their social networks), and measurable indicators that can be found outside the box that result from the contents of the black box. As a result, it is not necessary to gather data that directly answers the real question of interest. Since none of the data required must bear directly on the specific question at hand regarding the content of the black box, and since it may be easier to gather this seemingly innocuous information, the GM Meta Model can be used to reveal a true state in the absence of data. To meet this end, the GM Meta Model achieved the following:

- reconstructed the subjective decision processes of key leaders;
- incorporated the ways in which social structures and social interactions impact the decision preferences and choices of actors; and
- Enabled social network structures to influence pathway programs, potentially contrary to leadership goals or intentions, by creating social tension that affected the decision calculus of decision makers.

Additionally, the GM Meta Model semi-automated the data preparation process from other PAINT performer modules, SMEs, and social network models.

The GM Meta Model is a method to reveal information about a system through decisions made by the leadership; it supports simulation, forecasting, or reverse-engineering decision outcomes. Within the PAINT challenge problem the GM Meta Model was used to predict the decisions of individual leaders.

At the core of the GM Meta Model is the objective of implementing a generalizable decision model that has a verified theoretical basis as well as experimental results. Four social science theories were utilized to achieve these objectives, each of which NSI implemented and tested:

1. Expected utility decision making; for situations in which rational choice is possible (the classical approach to decision making);
2. Heuristic-based decision-making for situations of stress in which people do not behave “rationally” (using the heuristics of elimination-by-aspects, lexicographic, maximin heuristics that people revert to when making difficult decisions);
3. Group decision processes (aggregation of individual choices); and
4. Social influence’s role in weakening, strengthening, or changing the interests and attitudes of individuals (and thereby their decisions), as people’s interests dynamically respond to social interactions and environmental conditions.

The GM Meta Model included in the final demonstration at the PAINT PI meeting on January 8-9, 2009 in Reston, VA, included only a simple expected utility model, which contained social theories #1 and #3. The full GM Meta Model implements all four social theories, each of which has been validated.

Additionally, NSI aimed to and succeeded in designing a decision model that could be used to both 1) gain a deep understanding of the dynamics of a single key decision and 2) could model less specified generic decision-making. The final version of the GM Meta Model completed under this contract can be used to model any decision by any group of decision-makers and can scale to model a generic decision-making process.

3. METHODS, ASSUMPTIONS, AND PROCEDURES

To meet the PAINT integration goals, the Multi-Attribute Decision Model (MADM) module allows “generalized” decision-making that scales to any level of information in any domain (unlike conventional decision models) and reduces the data load concerning individual perceptions of potential outcomes and rankings of alternatives as much as possible through simplifying assumptions, discussed below. It also expands the decision-making calculations from expected utility to involve broader framing, heuristics, and the effects of social pressure. The model applies to any decisions we have data or prior history for, and the data/history itself does not need to be complete – the model will produce reasonable predictions about behavior, although without parameterization/error bars, since it has not **been tuned** for a particular group.

Although the current model is complex, it is also still a prototype. It displays the look and feel of human behavior, but was not developed on a particular set of leaders or particular decisions that have bearing on reality. As a result, the values in the model have not yet been calibrated to map precisely to real human behaviors or leadership styles. Although we would have liked to develop the model from the bottom up, studying numerous example decisions and validating at each step, we worked instead on the generalizable decision model to accommodate the PAINT demonstration schedule that needed inputs from us, before the more sophisticated versions of the decision model was completed.

The methods, assumptions and procedures used by NSI as part of the PAINT process can be classified into two areas: those that were used during research and development, and those that were used for collaboration and integration to a final model. This section is first divided into those two areas, and then the methods, assumptions and procedures of each are addressed. We begin by addressing the methods, assumptions, and procedures relating to the research and development of a generalizable decision model, and then discuss those relating to the collaboration and demonstration requirements of the PAINT model.

3.1 RESEARCH AND DEVELOPMENT

3.1.1 Methods

The GM Meta Model and the MADM module were implemented in Java code to enable easy integration with the other performers through the BAE backplane.

The GM Meta Model consists of five main software modules that integrate with the rest of the PAINT system: (i) a game- and decision-theoretic leadership *Multi-Attribute Decision Model (MADM)* to reconstruct the subjective decision problems of key leaders or leader types, which includes a *Socio-Cognitive Dissonance Module (SCDM)* to navigate the influence of peers and others in one's social network on one's values, preferences, and therefore decisions; (ii) a *Prognostic Bayesian Evaluator Model (PROBE)* to (a) explore the effects of varying decision task factors (e.g., problem complexity, time stress, choice rules) and decision makers' interests and preferences as represented in the MADM, and (b) incorporate a Bayesian evaluator to identify key intersections of other PAINT social network and developmental pathway models and areas demanding further investigation (e.g., probing to assess intent); (iii) a *Multi-Source Data Integrator Module (MSDIM)* which is a collection of social science tools and methods including social network analysis techniques to operationalize socio-cultural, process and historical data for input into the MADM component; and (iv) a *Scoring and Disambiguating Module (SADM)* which is a suite of statistical and related procedures to identify and rank the degree to which inputs into the decision model influence the decision outcomes. See Figure 1 on page 1 for a diagram.

The main focus of our nine month efforts was on the MADM module, due largely to demonstration collaboration requirements. However, in addition to completing the MADM, NSI has made significant progress toward completing the other aspects of the GM Meta Model ahead of schedule:

- *Multi-Source Data Integrator Module (MSDIM)* module for inputs includes work toward an SME interface to help SMEs who are not decision theory experts to inform the model. External domain knowledge is required to instantiate but not to run the model, and good information about decision-making context and decision maker preferences can be obtained from data analysis or SME inputs.
- *Prognostic Bayesian Evaluator (PROBE)* model as an overlay to the MADM module, which enables parameterization on outputs and sensitivity analysis on assumptions to more robustly characterize the decision space and situation at any point in time, has been initiated.

Note the following sections focus only on the MADM module.

3.1.2 Assumptions

The MADM module uses a dynamic approach to modeling the decision-making of individual leaders (actors) that emerges at confluence of four social theory assumptions:

- actors have preferences on which they evaluate options;
- actors are influenced by social networks and peers;
- actors adapt their responses to environmental conditions; and
- Actors contribute to a group that makes the final decision.

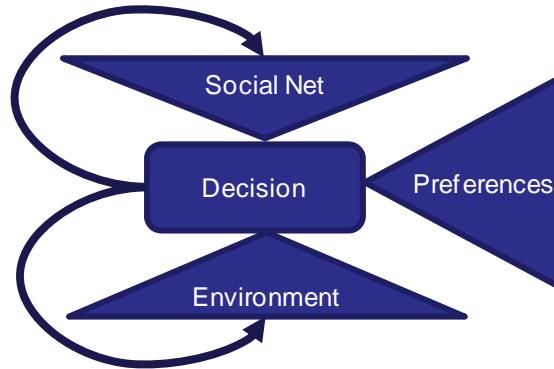


Figure 3: The MADM module assumes that decisions are affected by: 1) the social network, 2) the environmental context, and 3) the preferences of a decision maker

The first three assumptions combine to influence a single actor's decision. Each actor takes into account his/her own preferences, the direction in which s/he is pulled/pushed by his social network, and the environmental conditions in which the decision must be performed. His/her decisions in turn feed back into the surrounding social network and affect the decision environment of the future. In addition original preferences are updated as actors move in and out of new roles, which is not a continuous process (see Figure 3).

Modeling individual decision making is not sufficient to capture the leadership decision processes however, since most decisions are not made by individuals, but rather by groups that come together to discuss and finalize decisions. These groups are made up of individuals, each with preferences that influence the outcome, but the collective group dynamic adds additional complexity to the decision processes. The MADM module addresses these as well (see Figure 4).

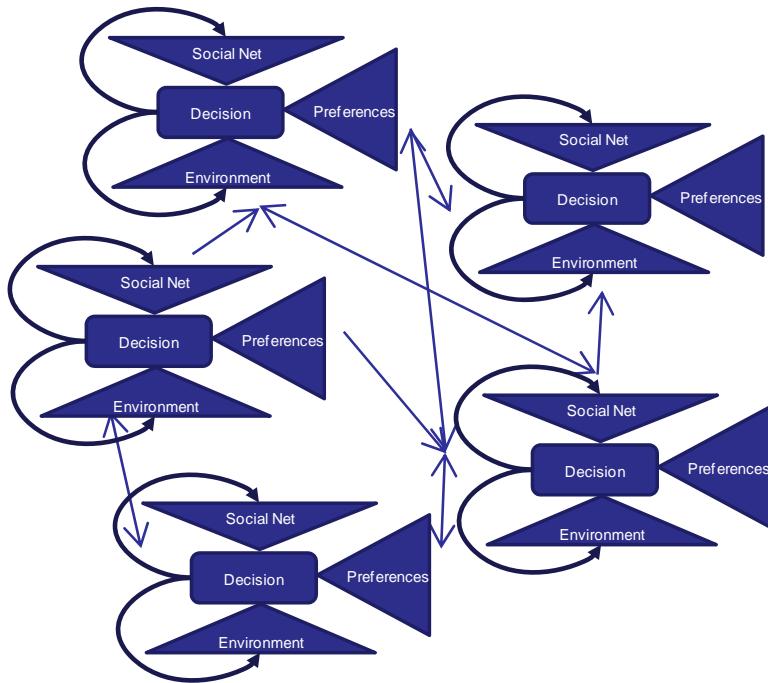


Figure 4: Decisions are not made in isolation; aggregation of group opinions must be addressed as well

The rest of this section will address how each of the assumptions was formalized into procedures.

3.1.3 Procedures

The MADM module addresses numerous topics as it progresses through the modeling of the decision process.

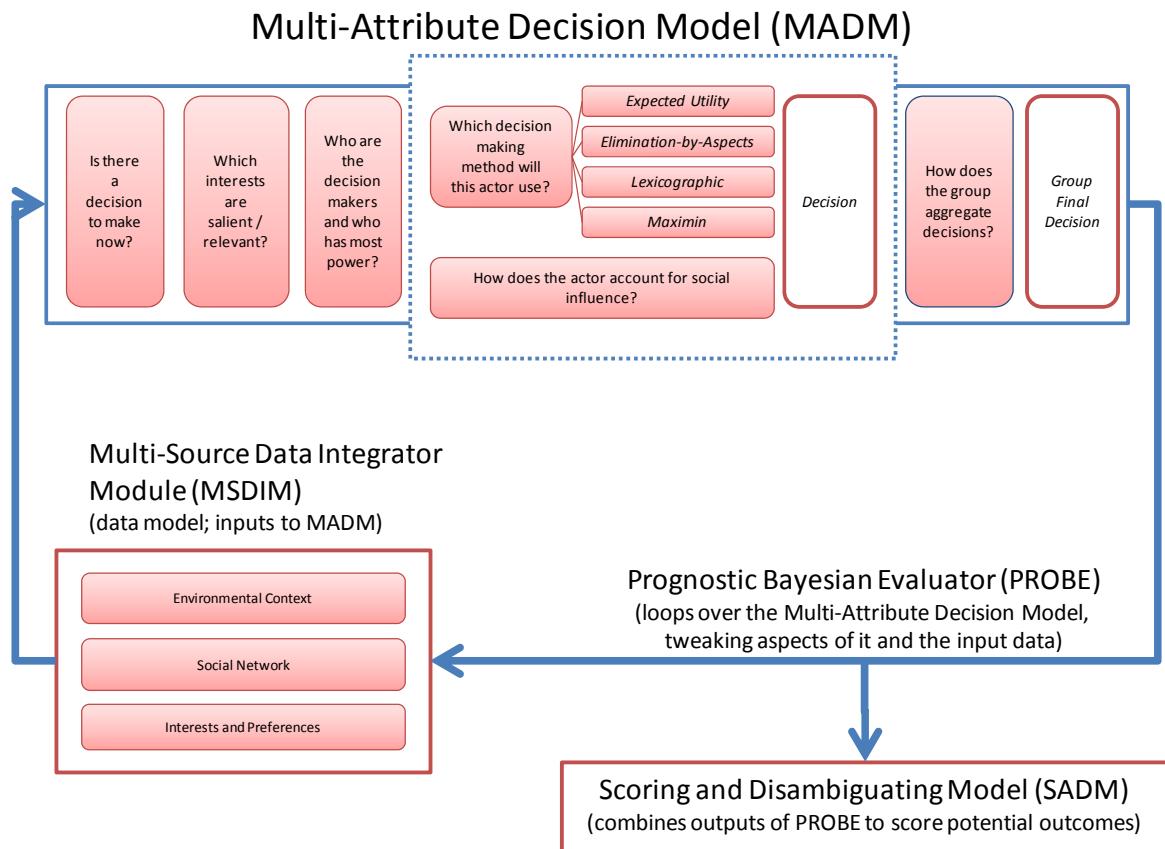


Figure 5: The MADM module addresses a number of questions in turn through its process and method of determining potential final decisions

This section addresses each aspect of the MADM module (see Figure 5). First we will discuss the issue of whether there is a decision to make, and which individuals and interests are most salient for that decision. Then we will discuss the method of making a decision for an individual – that is, the decision matrix model. Following that we will discuss, again on the individual actor level, how that actor accounts for social influence from his or her peers and superiors. Finally, after having calculated a set of individual decisions and the power relationships between them, the model aggregates the decisions to come to a final group decision. (These final decisions might then be run through the Prognostic Bayesian Evaluator, the results of which can be sent to the Scoring and Disambiguating Model to output to the rest of PAINT system. However, only the MADM module was completed during this effort and therefore only the MADM module will be discussed below.)

3.1.3.1 Is there a decision to make now? Which interests are salient/relevant? Who are the decision makers? Who has most power? - Environmental Conditions

The decision environment is dynamic, and its state at any point in time provides the decision that an actor will make. After identifying which decision needs to be made, the environmental context can mean that actors care about different things in making that decision – or the same things as before to a different extent than last week. Context also defines who the decision makers are for that particular decision, and the social network contains information both about who has the most formal power in that group, and who has the most informal power in that group – that is, influencing factors that will affect the final decision made.

3.1.3.2 Which decision-making model will an actor use? - Decision Matrices

The basis for the approach to decision-making taken in the MADM module is the *decision matrix*, a formalization for a single individual on which s/he performs decision calculus. The decision matrix includes his/her interests or decision dimensions (that is, the motivating factors), as well as the discrete alternatives s/he perceives as available to him/her (essentially the potential outcomes), and a “goodness” weight for each discrete alternative according to each interest (Reference [29]).

Figures 6-10 illustrate a sample decision process for an adversary's program development decision.

Adversary Program Development

Alternatives:		Decision Dimensions						Preferences	
		efficiency	marketability	schedule	profit	ethics	satisfaction	TOTAL	Overall rank
1A	Cut – money	4	4	7.5	7	7.5	6.5	36.5	8
1B	Cut – people	2	2.5	2.5	3.5	7.5	6.5	24.5	3
1C	Cut – time	3	2.5	5.5	5.5	4	6.5	27	5
1D	Cut – cost	1	1	2.5	1.5	2	6.5	14.5	1
2A	No cut - money	5	5	7.5	8	6	2.5	34	7
2B	No cut – people	6	6	2.5	3.5	5	2.5	25.5	4
2C	No cut – time	7	7	5.5	5.5	3	2.5	30.5	6
2D	No cut – cost	8	8	2.5	1.5	1	2.5	23.5	2

8=best ... 1=worst

↑
 Ordinal intra-dimension ranks;
 outcome value per dimension x
 alternative

Figure 6: Decision matrices formalize the decision-making process

The interests held by that leader that may influence this decision are listed across the top of the matrix ("efficiency," "marketability," "schedule," etc.). The perceived alternatives of that leader are listed in the first column. Each permutation of the alternatives is listed explicitly so it is equally formalized into the decision model. Finally, weights are assigned to each of the alternatives within each interest.

Once a decision matrix is formed, it becomes possible to perform calculations upon it. The most common calculation for this task is a simple expected utility calculation, which models perfect understanding of the decision environment and choices available, as well as perfectly rational actors who always maximize their interests and can perfectly determine the results of value trade-offs.

The same decision matrix shown in Figure 6 is pictured below in Figure 7, with the results of an expected utility/rational choice decision.

Adversary Program Development

Decision Dimensions

Alternatives:							TOTAL	Preferences Overall rank
	efficiency	marketability	schedule	profit	ethics	satisfaction		
1A Cut – money	4	4	7.5	7	7.5	6.5	36.5	8
1B Cut – people	2	2.5	2.5	3.5	7.5	6.5	24.5	3
1C Cut – time	3	2.5	5.5	5.5	4	6.5	27	5
1D Cut – cost	1	1	2.5	1.5	2	6.5	14.5	1
2A No cut – money	5	5	7.5	8	6	2.5	34	7
2B No cut – people	6	6	2.5	3.5	5	2.5	25.5	4
2C No cut – time	7	7	5.5	5.5	3	2.5	30.5	6
2D No cut – cost	8	8	2.5	1.5	1	2.5	23.5	2

rationalist choice rule

←

8=best ... 1=worst

Figure 7: Decision matrices can be operated upon using an expected utility/rational choice rule.

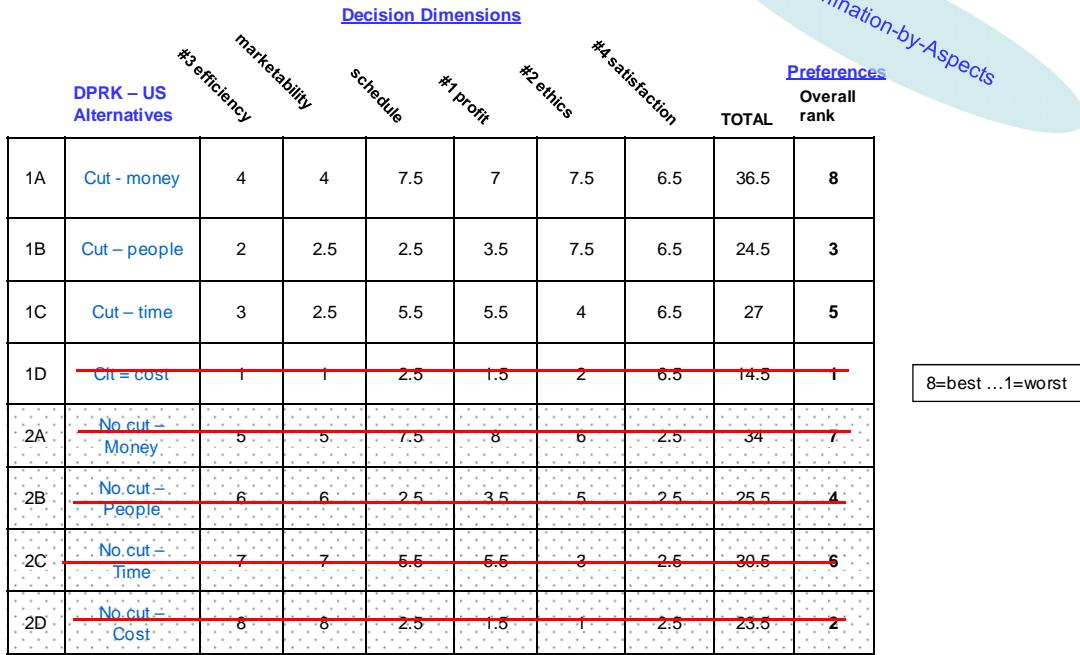
By adding across each row of the decision matrix, we can calculate the expected utility of each alternative. The alternative with the highest utility in the sample decision is for a cut in the amount of money provided (Reference [9]).

We know that individuals possess multiple interests, ambitions and objectives that condition their decision calculi and choice. However, research on decision processing (i.e., cognitive decision models) has increasingly shown that rather than a static tendency as once assumed, people in fact employ a range of decision strategies and rules (e.g., normative, bounded rationality, non-compensatory) that are contingent on the characteristics of the decision task (References [13],[20],[21]). Application of various choice heuristics to the “reconstructed” decision problems can help illuminate the range of likely choices – and most importantly, the factors that drive them – where essential choices must be made.

As a result, humans do not always act perfectly rationally. Therefore, the MADM module also incorporates multiple statistics for situations of stress in which an actor might not maximize his/her interests. In particular, the MADM module currently incorporates the elimination-by-aspects, lexicographic, and maximin decision-making heuristics. In this way, it extends and complements existing decision models.

The elimination-by-aspects heuristic eliminates the options with undesirable characteristics in turn (Reference [25]). In the below example, the four prioritized dimensions are efficiency, profit, ethics, and satisfaction. Finding the lowest value on four prioritized decision dimensions (columns), results in the elimination of those options; a tied lowest value results in the elimination of both of those options. Presuming all outcomes are believed to be credible, this process leads to the following decision matrix calculus:

Adversary Program Development



Decision Dimensions

DPRK – US Alternatives

Preferences
Overall rank

Elimination-by-Aspects

8=best ... 1=worst

		#3 efficiency	marketability	schedule	#1 profit	#2 ethics	#4 satisfaction	TOTAL	
1A	Cut - money	4	4	7.5	7	7.5	6.5	36.5	8
1B	Cut – people	2	2.5	2.5	3.5	7.5	6.5	24.5	3
1C	Cut – time	3	2.5	5.5	5.5	4	6.5	27	5
1D	Cut = cost	1	1	2.5	1.5	2	6.5	14.5	1
2A	No cut – Money	5	5	7.5	8	6	2.5	34	7
2B	No cut – People	6	6	2.5	3.5	5	2.5	25.5	4
2C	No cut – Time	7	7	5.5	5.5	3	2.5	30.5	6
2D	No cut – Cost	8	8	2.5	1.5	1	2.5	23.5	2

Figure 8: Decision matrices can be operated upon using an elimination-by-aspects heuristic.

The best option after an elimination-by-aspects calculation is to cut something, regardless of what. Elimination-by-aspects calculations work best for situations in which the options are represented by a large set of characteristics, and in which there are a large number of options – especially those in which a “first step” is needed to pare down the options to make a more nuanced decision reasonable.

The lexicographic heuristic compares alternatives along a set of prioritized dimensions (Reference [10]). The decision calculus for the lexicographic choice rule follows:

Adversary Program Development

Lexicographic choice rule

Alternatives:		Decision Dimensions						TOTAL	Preferences Overall rank
		#3 efficiency	marketability	schedule	#1 profit	#2 ethics	#4 satisfaction		
1A	Cut – money	4	4	7.5	7	7.5	6.5	36.5	8
1B	Cut – people	2	2.5	2.5	3.5	7.5	6.5	24.5	3
1C	Cut – time	3	2.5	5.5	5.5	4	6.5	27	5
1D	Cut – cost	1	1	2.5	1.5	2	6.5	14.5	1
2A	No cut – money	5	5	7.5	8	6	2.5	34	7
2B	No cut – people	6	6	2.5	3.5	5	2.5	25.5	4
2C	No cut – time	7	7	5.5	5.5	3	2.5	30.5	6
2D	No cut – cost	8	8	2.5	1.5	1	2.5	23.5	2

8=best ... 1=worst

Figure 9: Decision matrices can be operated upon using a lexicographic choice heuristic

The best option under this heuristic is *not* to cut, as the first priority is “regime survival” and the best alternative on that interest/decision dimension tells North Korea not to test. A lexicographic approach is intuitive and simple to understand, but it requires a non-arbitrary and transparent ordering of dimensions, which can be hard to perceive.

Finally, the MADM module also implements the maximin heuristic. The maximin calculation chooses the alternative that maximizes the minimum utility. That is, each alternative has a minimum amount to which it will be useful; the maximin calculation chooses the best of these, thereby minimizing its losses. The below decision matrix illustrates the maximin calculation, again presuming that all options are equally likely:

Adversary Program Development

Decision Dimensions

Alternatives:		Decision Dimensions								Preferences	
		#3 efficiency	marketability	schedule	#1 profit	#2 ethics	#4 satisfaction	TOTAL	Overall rank	Maximin Rank	
1A	Cut – money	4	4	7.5	7	7.5	6.5	36.5	8	8	
1B	Cut – people	2	2.5	2.5	3.5	7.5	6.5	24.5	3	3	
1C	Cut – time	3	2.5	5.5	5.5	4	6.5	27	5	5.5	
1D	Cut – cost	1	1	2.5	1.5	2	6.5	14.5	1	1.5	
2A	No cut – money	5	5	7.5	8	6	2.5	34	7	5.5	
2B	No cut – people	6	6	2.5	3.5	5	2.5	25.5	4	5.5	
2C	No cut – time	7	7	5.5	5.5	3	2.5	30.5	6	5.5	
2D	No cut – cost	8	8	2.5	1.5	1	2.5	23.5	2	1.5	

maximin choice rule

8=best ... 1=worst

Figure 10: Decision matrices can be operated upon using a maximin choice heuristic

The maximin calculation considers the worst consequence of each possible course of action, and then chooses the option that is the least worse – as a result, it models very cautious decision-making on the part of the actor (Reference [22]). Additionally, the maximin calculation does not allow for value trade-offs. In the program development example, the most cautious decision mainly have to do with cutting money or time allotted.

Formalizing a decision matrix that takes preferences of leaders into account is thus a very powerful technique that allows insight into how leaders might respond to particular decisions given different formulations of reality. The addition of heuristic decision-making modeling on top of the decision matrix framework enables an additional level of nuance in determining how the leaders are conceptualizing their world and what potential outcomes to expect with a given set of personal preference inputs.

3.1.3.3 How does the actor account for social influence? – Socio-Cognitive Dissonance

Human social interaction, decision-making and ultimately behavior is fundamentally constrained as well as enabled by the ties that exist between people (Reference [5]). Social structures -- religious, economic, political or societal -- impact the attitudes and behaviors of individuals by constraining and facilitating opportunities and options. Social network analysis provides a method for collecting, recording and analyzing empirical evidence of the ties between individuals. It also affords a formal, quantitative way to measure latent concepts such as culture, cohesiveness, power and influence which are critical for understanding group decision processes. Network groups and sub-groups can be

partitioned and assessed with social network methods such as N-cliques and block models. Group cohesiveness (tightness of connections) also can be measured (References [28], [30]), and key individuals can be discovered, e.g., with measures of centrality, prestige and power (Reference [28]). Networks, of course, are not static. Just as networks limit individuals, individuals and courses of actions affect structures like networks. The output of the MADM module will be used as input for other PAINT models. Moreover, despite leadership goals or intentions, social structures can interfere with the execution of programs and effect and redirect pathways. Clearly, effective modeling of any social system requires incorporation of these networks. This information is important in accurately defining the decision environment, and its inclusion in the MADM module extends and complements the decision model state-of-the-art.

The Socio-Cognitive Dissonance Module (SCDM) is designed to simulate the impact that social influence has on the actors' underlying cognitive structures. Given the network of social influence relationships among actors and the actors' cognitive structures at time t , the social influence model calculates the updated values of the actors' cognitive structures as they subconsciously, dynamically adapt their cognitive structures in response to their social influences (References [8], [14], [6]).

The first theoretical premise of the SCDM is that positive and negative influence in the social network translates into attraction and repulsion, respectively, in the cognitive space. Consider the relationship between an ego actor e and another (alter) actor a in the social network. If a has a positive influence on e , then e will display a tendency to adapt his/her cognitive structures to increase his/her cognitive similarity to a . However, if a has a negative influence on e , then e will display a tendency to adapt his/her cognitive structures to increase his/her cognitive differentiation from a . The result is that e is continuously exposed to multiple forces of cognitive attraction and repulsion to other actors in the network.

The second premise of the SCDM is that a discrepancy between an actor's social influence network and their cognitive position relative to other actors creates cognitive dissonance. In other words, if an ego e is positively influenced by an alter a , then high cognitive distance from (or dissimilarity to) a creates cognitive dissonance for e . Conversely, if an ego e is negatively influenced by an alter a , then high cognitive closeness (or similarity to) a creates cognitive dissonance for e .

The third premise of the SCDM is that making a change to one's cognitive structure also creates cognitive dissonance. In other words, the actors in the model exhibit cognitive inertia – a tendency to be predisposed toward propagating their existing cognitive structures.

The fourth premise of the SCDM is that each actor resolves these multiple forces by adapting his or her cognitive structure as needed to minimize his or her overall cognitive dissonance. This is posited to be a subtle, subconscious process, rather than an overt rational act.

We define a basic model of the cognitive dissonance experienced by the ego actor e as a result of the influence of the alter actor a as follows. Given two vectors of cognitive interests $\mathbf{w}, \mathbf{x} \in [0, 1]^m$, let $D(\mathbf{w}, \mathbf{x})$ denote the Euclidean distance between them (Equation 1).

Equation 1: Cognitive distance.

$$D(\mathbf{w}, \mathbf{x}) = \sqrt{\sum_{i=1}^m (w_i - x_i)^2}$$

Given an ego actor e and an alter actor a , let \mathbf{w}_a denote the cognitive interest vector of the alter a , and let η denote the social influence that alter a exerts over ego e . Then if ego assumes a given cognitive position ω , the cognitive dissonance incurred due to the influence of the alter is defined as:

Equation 2: Cognitive dissonance experienced by ego due to the influence of a single alter.

$$G(\omega | \mathbf{w}_a, \eta) = |\eta| * D(\omega, \mathbf{w}_a)^{2 * \text{sign}[\eta]}$$

Note that Equation 2 contains straightforward adaptations to the Coulomb's Law model as needed to satisfy the requirements of Attraction, Repulsion, Symmetry and Sign for the sociocognitive influence model.

The total cognitive dissonance experienced by ego e , including the cognitive dissonance arising from moving away from his/her own position, is calculated as the sum of the cognitive dissonance vis a vis ego him/herself and every other alter in the network. Let \mathbf{W} denote the matrix of the cognitive interests of all actors in the network that have a nonzero influence on e (including e him/herself), and let \mathbf{w}_j denote the vector of interests of actor a_j . Further, let η denote the vector of influences on the ego actor e (including e 's influence on him/herself), and let η_j denote the influence of actor a_j on the ego actor e . Then for a given cognitive position ω , the total cognitive dissonance that the ego e incurs by adopting the interests specified by ω is:

Equation 3: Total cognitive dissonance experienced by ego.

$$G(\omega | \mathbf{W}, \eta) = \sum_{j=1}^p |\eta_j| * D(\omega, \mathbf{w}_j)^{2 * \text{sign}[\eta_j]}$$

Thus the model of determining the new set of cognitive interests which the ego actor e will adopt in response to the social influence network is formulated as an optimization problem on the cognitive dissonance function. Let ω^* denote the new position of the ego actor e ; this position is calculated by solving the minimization problem:

Equation 4: New cognitive position of ego as a minimization of cognitive dissonance.

$$\omega^* = \operatorname{argmin} G(\omega | W, \eta) = \operatorname{argmin} \sum_{j=1}^p |\eta_j| * D(\omega, w_j)^{2 + \operatorname{sign}[\eta_j]}$$

The overall solution matrix Ω for the entire actor set is computed by calculating each actor's optimal new position independently and appending the actor's optimal interest vector to the matrix. Thus the row ω_j in the solution matrix Ω corresponds to the optimal new position ω_j^* for actor a_j computed according to Equation 4.

Each interest held by the actor represents a distinct dimension on which equilibrium must be found; as a result, at each time step, ego is being pulled and pushed in many different directions toward a final resting place at which the actor can find equilibrium. Each time step, the actor moves closer to what would be the state of equilibrium for him/her at that point in time. As actors move in and out of roles and positions around ego, and as ego's social network changes, however, the point of equilibrium toward which ego is making its way may also move.

For a visual representation of this process on a single interest, see Figure 11. The black line represents the direction in which ego will move this time step – that is, toward the minimization of cognitive dissonance that ego can find at about a weight of 0.5 on this particular interest. Ego is being pulled toward the horizontal positions of the alters in the positive space above the x-axis, and pushed away from the negative positions of the alters in the negative space below the x-axis. The further away from the x-axis that an alter is represented, the more ego is affected by that alter (positive attraction or negative repulsion.)

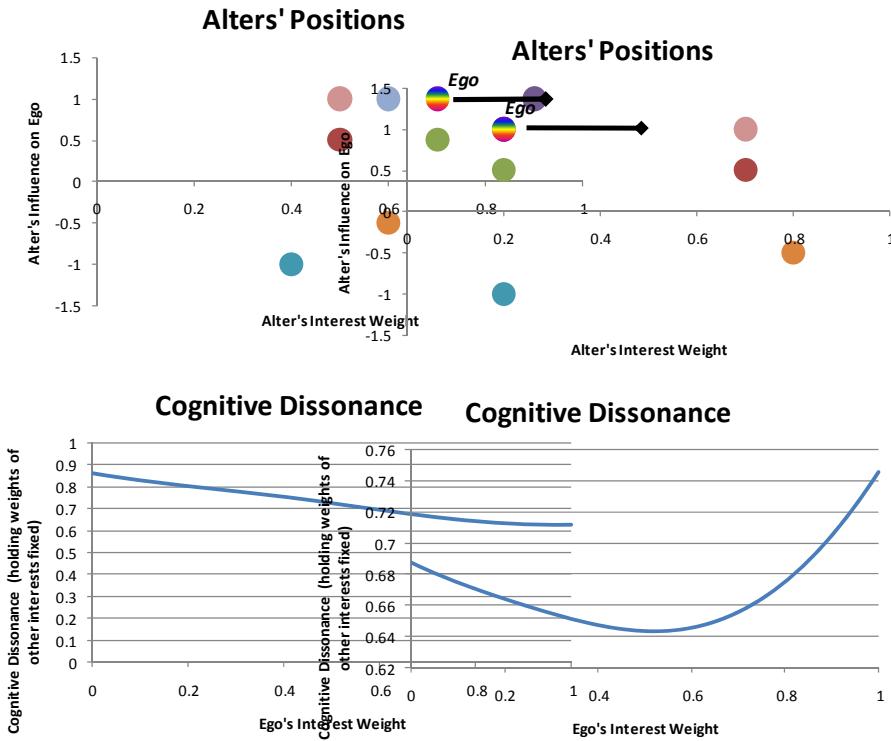


Figure 11: Social influence from peers and superiors pushes and pulls actors into new opinions and therefore new decisions.

Through the SCDM, an actor is influenced by his/her peers in terms of preferences and interests. The influence, of course, may not always or immediately be sufficient to invoke a change in decision outcome for that agent, especially as agents tend not to consciously re-make decisions they have previously addressed but rather rely on previous decisions to guide similar current ones.

3.1.3.4 How does the group aggregate decisions? – Group Decision Dynamics

In addition to individual or relatively cohesive group decision modeling, the MADM module also includes a facility to assess group preference aggregations and possible decision outcomes in instances where critical decisions are known to be made by decision units comprised of competing blocs and agendas (e.g., bureaucratic cabinets; consortiums of interest groups, etc.) Integration of this type of facility into the MADM module provides PAINT with a method for gaining insight into, and anticipating preference aggregation in, multi-party decision contexts.

Additionally, each decision leader has a decision-making style that s/he uses in groups s/he leads, following the Vroom & Yetton (1973) (Reference [27]) model of leadership and decision making. The three PAINT decision-making style categories are based on the basic categories of the Vroom-Yetton normative decision-making model (see Table 1 below). (The PAINT model does not distinguish the 8 subcategories.)

Table 1: The Vroom-Yetton model differentiates three main types of decision makers: autocratic, consultative, and group.

Group Leader Style	Decision-Making Process Styles	Description of small group decision-making process sequence and roles
Autocratic	A1- Autocratic or Directive	Leader defines issues, diagnoses problems, generates, evaluates and chooses alternatives
	A2- Autocratic with Group Input	Leader defines and diagnoses the problem; elicits inputs from group members
	A3- Autocratic with Group review and Feedback	Leader defines and diagnoses the problem; elicits inputs from group members: leader chooses solution and elicits feedback from group
Consultative	C1 –Individual Consultative	Leader defines problem; elicits ideas regarding the problem and solution alternatives. Leader chooses among alternatives.
	C2- Group Consultative	Same as C1 but leader shares problem definition process with the group.
Group	G1- Group Decision	Leader generates and shares the problem definition; group performs diagnosis, alternatives generation and evaluation; group chooses among alternatives.
	G-2 – Participative	Leader guides (facilitates) the group as a whole through the entire process
	G-3 Leaderless Team	The group has no designated leader; process leader or facilitator emerges.

These two techniques ensure that the MADM module captures the variety of group dynamics that exist in human decision-making populations, and that the final decision holds more closely to reality than a simple aggregation of individual choices.

3.2 COLLABORATION AND DEMONSTRATIONS

3.2.1 Methods

NSI discussed data requirements with BAE and Booz Allen Hamilton, and programmed a wrapper classes in Java that connected the MADM module to the backplane.

3.2.2 Assumptions

Decision making within the NSI model is generalized to be the result of two main entities with other entities in support (see Figure 12):

- decisions, each of which has a set of decision outcome alternatives and a set of interests that agents will consider in making that decision, and
- agents, each of whom has a particular importance weighting on that set of interests that come into play for a particular decision.

The connection between agents and the decisions they make (and the results they choose) are mediated by interests, or the behavioral causes that someone wants to fulfill. Interests refer to the concerns and priorities of the actors, and instantiation of interests in this realm results in a ranking of what an agent finds most important. For any decision, on the other hand, only a subset of interests are relevant – for instance, a decision about where to go to eat likely has nothing to do with one's interest in staying on schedule for one's project at work.

“Decisions” refer to “decision opportunities”. They are comprised of particular reasons that would influence someone toward a particular outcome (interests, as described above), and particular outcomes (decision alternative outcomes). Alternatives refer to a finite set of outcomes that may result from the decision opportunity. These opportunities are also defined by the information about the current state of the system at the time of the decision. Finally, agents include decision-makers and all others who can influence the decision maker. In addition, information on the agents included how they were connected to each other and information on their roles, responsibilities, status within the hierarchy etc. Agents have interests that relate to the outcome alternatives.

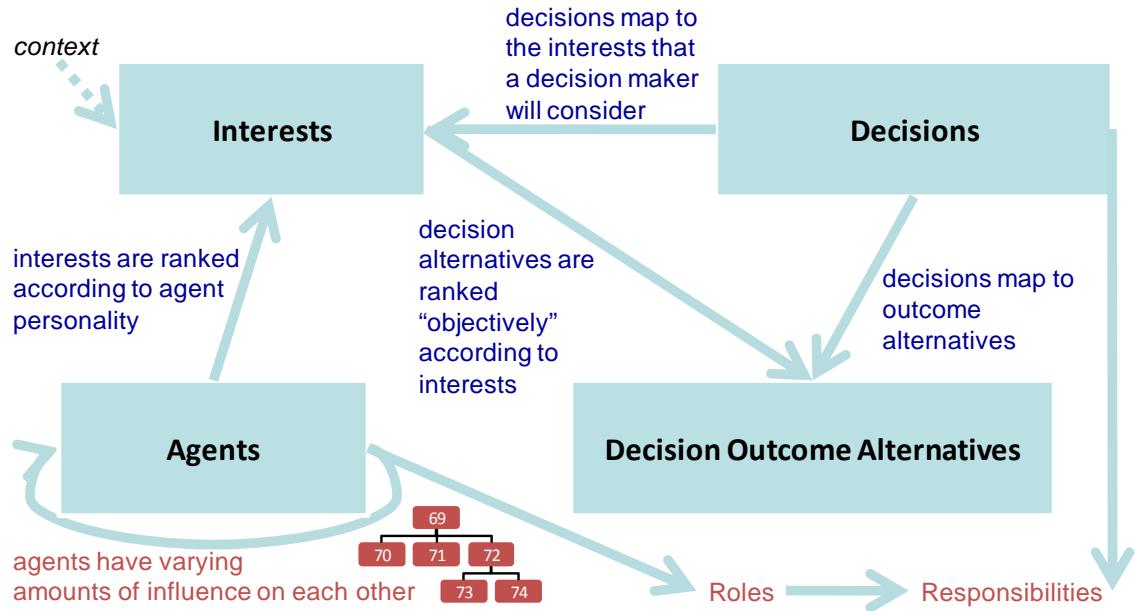


Figure 12: Accurate data for the decision model is assumed to come from many different performers

Additionally, not included in the chart above but parallel to it, are the roles that any agent may fill, which have the same relation to the other elements of the model that agents have. However, agents may move in and out of roles as needed and as promoted/demoted by their organization.

In order to reduce the data requirements for the demonstrations, which would otherwise require each agent to have a separate subjective decision matrix in which the alternatives were ranked according to the subjective perception of their “goodness” for a particular interest, NSI has instead made the simplifying assumption that for each decision there is an “objective” ranking of the alternatives according to the interests. Agents then differ only on the extent to which they weight each interest. This simplifying assumption makes the problem of a generalizable decision model tractable, as the design choice greatly reduces the number of inputs required to instantiate the model, from:

$\#decisions * \#interests * \#alternatives * \#agents$ (the decision matrices) +
 $\#agents * \#interests$ (the importance each agent puts on each interest)

to

$\#decisions * \#interests * \#alternatives$ (the single “objective” decision matrix) +
 $\#agents * \#interests$ (the importance each agent puts on each interest)

thereby reducing the amount of input data for the model by a factor of the number of agents. For the PAINT challenge problem, any agent in the Booz Allen Hamilton social network model could take on any decision-making role, so the reduction was significant by a factor of 107.

3.2.3 Procedures

The Multi-Source Data Integrator Module (MSDIM) takes in inputs from files, the backplane, and other models. At each time step, the MADM module outputs the alternative chosen by a group of decision makers for each decision included in the input data; it also exports the personal decisions of those involved in the decision at each time step. Additionally, the model can be probed through toggling the relevance of an interest “off” (and potentially later back “on”), to model the elimination of one consideration from the decision calculus of the agents in the model.

During instantiation, the MSDIM reads data from a set of input files external to the model, which form a relational database. To this end, the user is prompted to identify a .dmd (“decision matrix data”) file at the beginning of each run.

Each line in the .dmd file identifies the location of the corresponding input matrix (in comma-delimited format, with top row containing labels describing the rest of the data) required by the MADM module. The first element of the file identifies which matrix is being addressed on this line; a comma separates the pre-defined label for that matrix from its location. Table 2 illustrates what a .dmd file looks like.

Table 2: An example decision-making data file (.dmd file) format.

decisionMatrixFile, "H2Decisions.csv"
alternativesMatrixFile, "H2Alternatives.csv"
interestMatrixFile, "H2Interests.csv"
altToInterestMatrixFile, "H2Alternatives-Interests.csv"
agentPersonalityMatrixFile, "H2Agents-Interests.csv"
rolePersonalityMatrixFile, "H2Roles-Interests.csv"
decToInterestMatrixFile, "H2Decisions-Interests.csv"
agentsToStyleMatrixFile, "H2Agents-Style.csv"
contextCorrectionMatrixFile, "H2ContextCorrection.csv"

The data matrices of inputs must be listed and linked with comma-delimited files on separate lines within that .dmd file (here the matrix files are represented as .csv files, but that extension is not required). Table 3 illustrates how each matrix file must then have the expected content in a processing format.

Table 3: The MADM module accepts formalized input during instantiation that defines the decision makers, their preferences, and the decisions that may need to be made, in a particular format

Label in .dmd File	Content of .csv file	Internal Format of .csv File
decisionMatrixFile	Describes the decisions.	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>DecisionID</i> (int) • <i>Description</i> (String, short description) • Row 2 begins the data following this format
alternativesMatrixFile	Describes the alternatives.	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>DecisionID</i> (relational int; matches the decision to which this alternative belongs) • <i>AlternativeID</i> (int, starting at 0) • <i>Alternative Description</i> (String, short description) • Row 2 begins the data following this format
interestMatrixFile	Describes the interests.	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>InterestID</i> (int, starting at 1) • <i>Description</i> (String, short description) • Row 2 begins the data following this format
altToInterestMatrixFile	<p>Maps the alternatives to the interests.</p> <p><i>Conceptually:</i> Each interest has a ranking of alternatives that would best fulfill it; this matrix expresses that ranking.</p>	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>Decision</i> (relational int; matches the decision which these interests are describing) • <i>Alternative ID</i> (relational int; matches the alternative which these interests are describing) • each interest ID, in same order as was presented in interestMatrixFile (ints, where each alternative for a particular decision is ranked)

		<p>in the range [0, #alternativesForThisDecision], with the mean used for ties)</p> <ul style="list-style-type: none"> • Row 2 begins the data following this format
agentPersonalityMatrixFile	<p>Maps the agents to weights they put on interests.</p> <p><i>Conceptually:</i> Each agent cares about certain interests more than others and will persist in caring about these interests no matter which role that agent currently fills; this matrix expresses those preferences.</p>	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>Person ID</i> (relational int; matches with the agent identifiers provided to Booz Allen Hamilton's Social Network Model) • each interest, in same order as was presented in interestMatrixFile (ints, where each interest is weighted in the range [0, 10], with 10 meaning that interest has extremely high importance to that agent) • Row 2 begins the data following this format
rolePersonalityMatrixFile	<p>Maps the roles to weights they put on interests.</p> <p><i>Conceptually:</i> Each role has certain interests that anyone in that role is concerned with to a particular extent; this matrix expresses that "job description" information.</p>	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>Role ID</i> (relational int; matches with the role identifiers provided to Booz Allen Hamilton's Social Network Model) • each interest, in same order as was presented in interestMatrixFile (ints, where each interest is weighted in the range [0, 10], with 10 meaning that interest has extremely high importance to that agent) • Row 2 begins the data following this format
decToInterestMatrixFile	Maps the decisions to which interests are considered in making that	<ul style="list-style-type: none"> • Row 1 must include labels for <ul style="list-style-type: none"> • <i>DecisionID</i> (relational int; matches decision that is being described)

	decision, as all interests may not be used for all decisions.	<ul style="list-style-type: none"> each interest, in same order as was presented in interestMatrixFile (boolean, where each interest is 0/“irrelevant” or 1/“relevant” for the particular decision in that row) Row 2 begins the data following this format
agentsToStyleMatrixFile	Expresses the decision-making style of each agent (<i>autocratic</i> , <i>consultative</i> , or <i>group</i> , based on the Vroom-Yetton normative decision-making model).	<ul style="list-style-type: none"> Row 1 must include labels for <ul style="list-style-type: none"> <i>PersonID</i> (relational int; matches the agent who is being described) <i>Decision-making Style</i> (int, from set {1,2,3}; 1 indicates autocratic decision-making style, 2 indicates consultative decision-making style, 3 indicates group decision-making style) Row 2 begins the data following this format
contextCorrectionMatrixFile	Expresses the reaction of each role to suboptimal performance on each interest (based on a typology created by BAE: <i>resolution</i> , in which a leader increases interest, or <i>resignation</i> , in which a leader decreases interest).	<ul style="list-style-type: none"> Row 1 must include labels for <ul style="list-style-type: none"> <i>Role ID</i> (relational int; matches the role that is being described) each interest, in same order as was presented in interestMatrixFile (int, from set {1,2}; 1 indicates inclination of that role to resolution on that particular interest, 2 indicates inclination of that role to resignation on that particular interest) Row 2 begins the data following this format

For the purposes of the demonstrations, this initialization data was provided for the particular PAINT challenge problem by BAE. The file provided by NSI for integration that contains the MADM module and MSDIM also includes the input data created for this demonstration so that all information is packaged together. These matrix file inputs, however, can of course be modified at any point in time for an immediate effect.

During runtime, NSI also receives environmental context feedback about the extent to which each of the interests is currently being fulfilled in the other PAINT performer models. This feedback comes in the form of context variables that vary from 1 (optimal fulfillment of interest) to 0. The feedback variables expected on the backplane follow the following format:

INTEREST_NAME_Status

where `INTEREST_NAME` is the name of the interest as defined in the input files, with spaces replaced by underscores to accommodate the backplane.

For instance, the expected feedback variable for an interest labeled “V1 Schedule” in the `interestMatrixFile` would be `V1_Schedule_Status`. Similarly, the corresponding feedback variable for an interest labeled “Performance of V2” would be `Performance_of_V2_Status`. The context variables expected will vary with the interests defined at start-up.

Additionally, Booz Allen Hamilton provides important statistics about its social network to NSI at every time step, enabling the socio-cognitive dissonance calculations in which an agent gradually moves to align him/herself in a position on each interest that minimizes the socio-cognitive dissonance endured by that agent as a result of similarity to agents disliked by that agent and dissimilarity to agents liked by that agent. The socio-cognitive dissonance module was discussed earlier in the Technical Approach section.

At every time step, all the decisions that were defined in the input .csv files are run by the MADM module. A particular alternative is chosen for every decision by a group at every time step, which is output to the backplane to be routed into the appropriate models. A decision regarding the progress of pathways, for instance, would have an effect on the pathways in the next time step.

Furthermore, as requested by the BAE integration team, NSI implemented an action that allows a probe to “turn off” an interest for all decisions.

Table 4: The probe to “toggle off” an interest’s relevance to individual decision makers can be taken against the decision model

Action	Parameters	Description
<i>Toggle Interest Relevance</i>	enum: all defined interests enum: “off” “on”	By toggling the interest relevance, a probe can eliminate an interest from consideration in decision-making; the interest can also be reinstated.

By toggling interest relevance, a probe can stop an interest from being considered by the decision makers; that same interest can be turned back on at a later time step (see Table 4).

This action was requested by BAE to simulate the effect of a press release by a competing company that alerts the world to the fact that Company X is working on product V1; BAE expected that the interest “Proprietary of V1” would therefore no longer factor into the decision calculus of the organization. Finer granularity was not requested for this demo, but can be provided at request in the future.

Note that if an interest was marked as irrelevant to a particular decision *a priori*, toggling that interest will not enable it to be used in that decision (the `decToInterestMatrixFile` indicates which interests are relevant to which decisions). (For instance, the interest “having vegetarian options,” although relevant to a decision regarding a restaurant for lunch, is irrelevant to a decision regarding the amount of funding to allocate to R&D. Toggling the interest “having vegetarian options” is therefore always meaningless to a decision regarding funding allocation. As a result, toggling an interest will turn it off or on only for those decisions that already find that interest to be relevant.)

4. RESULTS AND DISCUSSION

Progress was made internally in the decision model that was formalized, as well as externally through using social network data provided by Booz Allen Hamilton and initialization data provided by BAE. A discussion of the results of the model running independently and the model when provided with Booz Allen Hamilton, BAE, and other performer data follows.

4.1 RESEARCH AND DEVELOPMENT

Despite a shortened timeframe, NSI has completed the MADM module of the GM Meta Model, including the SCDM, and has made progress toward all other proposed models except the Scoring and Disambiguating Module (SADM). The MADM module addresses leadership decision-making through integrating the environment, social network, and preferences of individuals, and then modeling the group decision-making processes of those individuals according to verified social science theory.

Additionally, we have shown that it is possible to reveal what is inside a black box by looking at the measurable indicators of the decisions made inside that box. Subtle changes internal to the box can affect outcomes, and it is possible to capture these changes. Additionally, some decision situations are more robust than others to change.

4.2 COLLABORATION AND DEMONSTRATIONS

Numerous experiments have been performed using the MADM module, both integrated into the other models and separately. Almost all experiments have been performed using the instantiation data provided by BAE for the PAINT challenge problem, regarding a product development company that has multiple product lines (V1, V2, V3) under development, and thus using collaboration data.

Given changes in the decision-making environment in the integrated model, actors' decisions change. Figure 13 offers one example in which in the integrated model, the V2 pathway (top center) was deliberately starved from the outside for funding and materials. In response, the decision regarding funding (bottom) changed from funding all the vaccine lines equally, to funding V2 at 50% and V1 and V3 at 25% each. When the pathway is starved, in other words, decision makers respond to fund it more extensively.

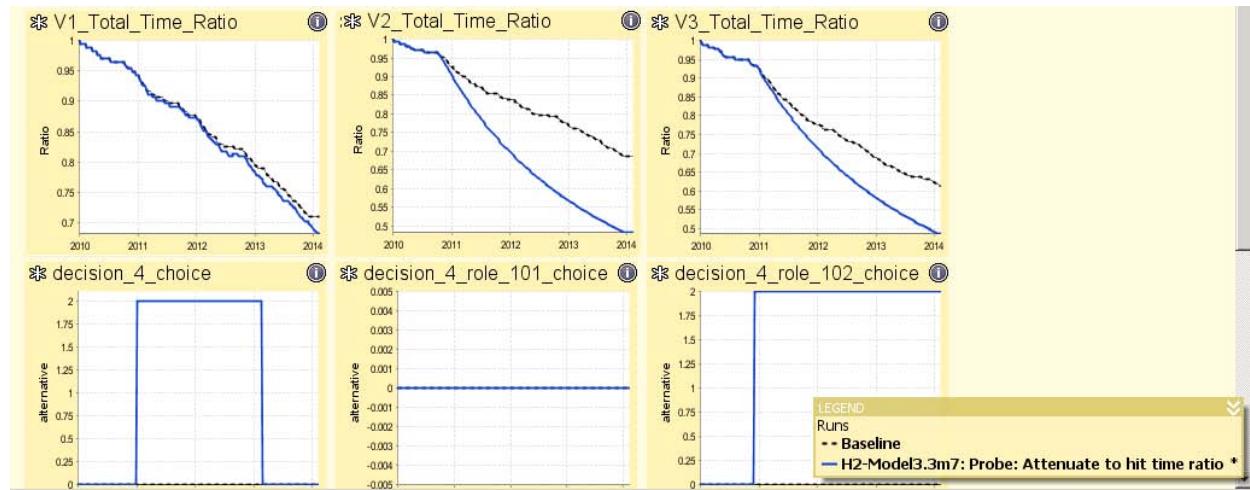


Figure 13: In the integrated model, when a pathway is starved, decision makers respond to fund it more extensively.¹

As the time to completion ratio for V2 drops, within a few weeks, the leadership reacts by changing their funding from allocating 33% to V2 to allocating 50% to V2. In particular, although the person in role 101 does not change his decision, the individuals in role 102 and role 105 now believe that 50% of the resources should go to V2 and therefore the final decision choice for this issue is to allocate 50% of funding to V2.

¹ In terms of the consequences of the change in decision to fund V2 with 50% of funds: V1 and V3 slipped schedule as well (top left and right). However, despite the intent decision to increase the funding to V2, no materials are available to improve the time ratio to completion of V2, so it stays behind schedule (top center).

Additionally, the SCDM has been shown in experiments using integrated and dynamic data from Booz Allen Hamilton to have an effect on the agents' decisions. For instance, one experiment showed that when accounting for social influence from others who have different priorities, one agent is inclined to fund all lines equally. However, if he is isolated from his social network, which favors other product lines, and is not subject to influence from them, he would rather fund V2 especially. We therefore see an experimental difference between the states of the model if the influence of agents on each other is enabled versus when it is not.

Similar results have been found with other experimental tests. For instance, initially one agent in a particular role favors funding V3. At each time step, however, he accounts for the influence of his alters upon him – and the alters that he dislikes are part of the faction that prefers V3, whereas the actors that he likes are part of the faction preferring to fund V1 and V2. The likelihood of this agent funding V1 grows more substantially once the alters are accounted for. That is, the agent is being pulled and pushed by his social network to favor V1 over V3, and does so – although he is part of a faction that *a priori* prefers funding V3. This experimental outcome is both reasonable and provides insights into the social dynamics of decision-making that otherwise may be overlooked.

Additionally, integration of the decision model with the social network model has other interesting and even perhaps counter-intuitive results that nevertheless are what reasonable humans may do. For instance, we have found in experiments that once a decision maker for funding issues is isolated, decision makers may defer to someone new in making decisions. This occurs because an isolated decision maker means changes in levels of informal power in the organization, so deference within a decision-making group may go to someone else. Similarly, we have found that isolation of an agent in one role may lead to another role changing his decision – this is because the second actor is influenced by the first actor; without having the first actor looking over his shoulder, the second actor is free to “act his conscience” and vote to fund product lines or allocate other resources as he wishes.

The social relations of humans underlie much of group decision making, and a change in the social network may have diverse and rippling effects on the decision-making environment. For instance, social network changes can affect the extent to which an individual perceives an interest as “important,” what the individual’s peers’ interests are, and who is judged “most powerful” and therefore has most influence on the group. Decisions change based on social ties, and isolation of individuals can change power dynamics and therefore decisions. Social structures can interfere with programs, and affect and redirect pathways despite leadership goals or intentions. The behaviors that emerge from modeling of these phenomena are not apparent but nevertheless make sense with post-hoc analysis and seem very human.

From these results that indicate the very diverse effects that changes to social networks can have on decisions, it seems important to continue to integrate social science and social networks into decision models that formalize the decision processes of actors. This track of inquiry should persist into the future.

5. CONCLUSIONS

Despite a shortened timeframe, NSI has completed the MADM module of the GM Meta Model, including the SCDM, and has made progress toward all other proposed models. The MADM module models complex human decision making subject to a broad range of influences at the individual and group levels. MADM decision-makers respond in a variety of ways to a variety of inputs. Depending on external inputs, rational choice (expected utility) or a range of heuristic models are invoked to model the variance in human decision models. Actor's interests, preferences, role, personality, and social pressure are all factored into the decision calculus of each decision maker for each decision. Group dynamics are carefully modeled to ensure aggregate outcomes reflect processes for group decision making.

In addition to completing the MADM module, NSI has made significant progress toward completing the other aspects of the GM Meta Model ahead of schedule:

- Multi-Source Data Integrator module for inputs includes work toward an SME interface to help SMEs who are not decision theory experts to inform the model. External domain knowledge is required to instantiate but not to run the model, and good information about decision-making context and decision maker preferences can be obtained from data analysis or SME inputs.
- Prognostic Bayesian Evaluator module as an overlay to the Multi-Attribute Decision Model, which enables parameterization on outputs and sensitivity analysis on assumptions to more robustly characterize the decision space and situation at any point in time, has been initiated.

In the process of creating the GM Meta Model, NSI has created a new approach to formalizing an adaptable and generalizable decision model that can be applied to any domain where the options are discrete and the decision criteria are independent. Even a large number of decision makers can be included in this decision-making process. Additionally, where other decision models use purely an expected utility calculation or integrate with social networks simply through group bargaining models, the GM Meta Model extends beyond the extensive focus on only non-cooperative game theory to include the ability for non-expected utility decision choice heuristic options, as well as the effect of social influence and the influences from the environment and dynamic changes in role. Such work to such a scale to our knowledge has not previously been performed.

6. RECOMMENDATIONS

New and remaining open research issues in the social sciences that would be helpful to better determining the decision model include issues surrounding modulation of it. For instance, further research should be performed on the timing of cognitive change. How gradually do individuals align themselves according to their social networks?

Additionally, further research should be performed on the process of triggering heuristic use instead of an expected utility calculation. In particular, what contextual and environmental states influence someone in favor of triggering into heuristic use? Additionally, what sorts of personal and contextual indicators can be used to inform the likelihoods with which someone will use a particular heuristic?

Finally, the fidelity and validity of the MADM model depends to some degree on the accuracy and richness of the data on actor interests, priorities and personal relationships. In its conception, therefore the GM Meta Model included the Multi-Source Data Integrator, to provide an interface to input, clean and integrate different data types. Because data is often unavailable or unreliable the GM Meta Model also requires a means to measure model outcomes under uncertainty. The purpose of the Prognostic Bayesian Evaluator Module was to determine the robustness of decisions in the event that a given set of assumptions (inputs) about reality are incorrect, and to parameterize the degree to which a particular decision outcome is likely given a specific state of reality.

7. REFERENCES

- [1] Astorino-Courtois, Allison. (1996). "Transforming International Agreements into National Realities: Marketing Arab-Israeli Peace in Jordan," *The Journal of Politics*. 58(4), November, pp. 1035-1054.
- [2] Astorino-Courtois, Allison. (1998). "Clarifying Decisions: Assessing the Impact of Decision Structures on Foreign Policy Choices during the 1970 Jordanian Civil War," *International Studies Quarterly*, 42, pp. 733-754.
- [3] Astorino-Courtois, Allison. (2000). "Can peace be marketed? A preliminary analysis of Israelis and Palestinians," *Conflict Management and Peace Science*, 18(1).
- [4] Bettman, J. and Eric Johnson. (1993). *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.
- [5] Bueno de Mesquita, Bruce and Rose McDermott. (2004). "Crossing No Man's Land: Cooperation from the Trenches," *Political Psychology*, Vol. 25, No. 2: 271-287.
- [6] Cartwright, D., and Harary, F. (1979). Balance and clusterability: An overview. In Holland, P.W., and Leinhardt, S. (eds.), *Perspectives on Social Network Research*. New York: Academic Press, pp. 25-50.
- [7] Druckman, James. (2004). "Political Preference Formation: Competition, Deliberation, and the (Ir)relevance of Framing Effects," *The American Political Science Review*, Vol. 98, No. 4, pp. 671-686.
- [8] Festinger, L. (1962). *A Theory of Cognitive Dissonance*. London: Tavistock.
- [9] Fishburn, P. C. (1970). *Utility Theory for Decision Making*. New York: Wiley.
- [10] Fishburn, P. C. (1974). Lexicographic Orders, Utilities and Decision Rules: A Survey. *Management Science*, 20(11), 1442-1471.
- [11] Frisch, Deborah. (1993). "Reasons for Framing Effects," *Organizational Behavior and Human Decision Processes*, 54, pp. 339-429.
- [12] Geva, Nehemia and Alex Mintz. (1994). "Framing the Options for Peace in the Middle East," Conf. paper delivered at ECAAR-Israel Conference, Haifa, Israel, 20 June 1994.
- [13] Hansen, David and James Helgeson. (1996). "Choice Under Strict Uncertainty: Processes and Preferences," *Organizational Behavior and Human Decision Processes*, 66(2), May, pp. 153-164.

[14] Heider, F. (1979). "On balance and attribution." In Holldan, P.W., and Leinhardt, S. (eds.), *Perspectives on Social Network Research*. New York: Academic Press, pp. 11-23.

[15] Hudlicka, E., Karabaich, B., Pfautz, J., Jones, K., & Zacharias, G. (2004). Predicting Group Behavior From Profiles and Stereotypes. In *Proceedings of Behavior Representation in Modeling and Simulation (BRIMS)*.

[16] Kahneman, Daniel and Amos Tversky. (1979). "Prospect theory: An analysis of decision under risk," *Econometrica*, 47, pp. 263-291.

[17] Kahneman, Daniel and Amos Tversky. (1984). "Choices, Values and Frames," *American Psychologist*, 4, pp. 341-350.

[18] Maoz, Zeev. (1990). *National Choices and International Processes*. Cambridge: Cambridge University Press.

[19] Paletz, David and Robert Entman. (1981). *Media, Power, Politics*. NY: The Free Press.

[20] Payne, John W., James Bettman and Eric Johnson. (1993). *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.

[21] Payne, John, James Bettman and Eric Johnson. (1988). "Adaptive strategy selection in decision making," *Journal of Experimental Psychology: Learning, Memory and Cognition*. 14, pp. 534-552.

[22] Pearman, A. D. (1977). A Weighted Maximin and Maximax Approach to Multiple Criteria Decision Making. *Operational Research Quarterly* (1970-1977), 28(3), 584-587.

[23] Quattrone, George and Amos Tversky. (1988). "Contrasting Rational and Psychological Analyses of Political Choice," *American Political Science Review*. 83: 719-736.

[24] Russell, J. and Norvig, P. (2003). *Artificial Intelligence: A Modern Approach*. Pearson Education, Upper Saddle River, NJ.

[25] Tversky, A. (1972). "Elimination by aspects: a theory of choice." *Psychological Review*, Vol. 79: 281-299.

[26] Vroom, V. H. (2000). Leadership and the decision-making process. *Organizational Dynamics*, 28(4), 82-94.

- [27] Vroom, V. H. (1973). *Leadership and Decision-Making*. Pittsburgh: University of Pittsburgh Press.
- [28] Wasserman, Stanley, and Katherine Faust. (1994). *Social Network Analysis: Methods and Applications*. Cambridge: Cambridge University Press.
- [29] White, D. J. (1969). *Decision Theory*. Chicago: Aldine Pub. Co.
- [30] White, Douglas and Frank Harary. (2001). The Cohesiveness of Blocks in Social Networks: Node Connectivity and Conditional Density. *Sociological Methodology* 31:305-359.
- [31] Zelik, D., Patterson, E., & Woods, D. (2007). Judging Sufficiency: How Professional Intelligence Analysts Assess Analytical Rigor. In Proc. of Human Factors and Ergonomics Society 51st Ann. Mtg.

8. LIST(S) OF SYMBOLS, ABBREVIATIONS, AND ACRONYMS

- .csv – comma-separated value (file format)
- .dmd – decision model data (file format used to bring individual input matrix files into decision model)
- AFMC – Air Force Material Command
- AFRL – Air Force Research Laboratory
- BAE – BAE Systems (performer on PAINT; integrator role)
- CBRN – chemical, biological, radiological, nuclear
- DNI – Director of National Intelligence
- GM – Goals-Means
- IARPA – Intelligence Advanced Research Projects Activity
- MADM – Multi-Attribute Decision Model
- Maximin – decision rule: “maximize the minimum possible payoff” (in decision theory, game theory, statistics, philosophy)
- MSDIM – Multi-Source Data Integrator Module
- NSI – National Security Innovations, Inc.
- PAINT – Pro-Active INTelligence
- PROBE – Prognostic Bayesian Evaluator
- SADM – Scoring and Disambiguating Module
- SCDM - Socio-Cognitive Dissonance Module
- SME – subject matter expert
- US – United States
- V1; V2; V3 – product lines under development on PAINT Challenge Problem data provided by BAE; each product line has associated interests and plays into particular decisions